

Transport on Adaptive Random Lattices

Jelle Ritzerveld* and Vincent Icke
Sterrewacht Leiden, P.O. Box 9513, 2300 RA Leiden, The Netherlands
(Dated: February 6, 2008)

In this paper, we present a new method for the solution of those linear transport processes that may be described by a Master Equation, such as electron, neutron and photon transport, and more exotic variants thereof. We base our algorithm on a Markov process on a Voronoi-Delaunay grid, a nonperiodic lattice which is derived from a random point process that is chosen to optimally represent certain properties of the medium through which the transport occurs. Our grid is locally translation and rotation invariant in the mean. We illustrate our approach by means of a particular example, in which the expectation value of the length of a grid line corresponds to the local mean free path. In this example, the lattice is a direct representation of the ‘free path space’ of the medium. Subsequently, transport is defined as simply moving particles from one node to the next, interactions taking place at each point. We derive the statistical properties of such lattices, describe the limiting behavior, and show how interactions are incorporated as global coefficients. Two elementary linear transport problems are discussed: that of free ballistic transport, and the transport of particles through a scattering medium. We also mention a combination of these two. We discuss the efficiency of our method, showing that it is much faster than most other methods, because the operation count does *not* scale with the number of sources. We test our method by focusing on the transport of ionizing radiation through a static medium, and show that the computed results for the classical test case of an ionization front expanding in a homogeneous medium agree perfectly with the analytic solution. We finish by illustrating the efficiency and flexibility of our method with the results of a simulation of the reionization of the large scale structure of the Universe.

PACS numbers: 05.10. Gg, 51.10.+y, 95.30.Jx

I. INTRODUCTION

A. Grids and errors

The numerical description of physical systems that obey a differential expression requires the use of a differencing technique on some sort of grid, mesh or lattice (except in rare cases where computational symbolic algebra can be applied). Accordingly, the question is not if errors are made, but rather what kind, and of what severity. In the past, computational lattices were almost always regular and rectangular, the Cartesian grid being the archetypal example. Analytic estimates of the corresponding errors are routinely made for ordinary differential equations on such grids (e.g. [1]); in the case of partial differential equations, especially those of hydrodynamics, the error analysis may even be quite elegant (see e.g. [2], [3], [4]).

In almost all of these cases, the numerical method may be chosen in such a way that the errors become small with increasing grid resolution. In practice this is often less than helpful, due to the steep increase of computational effort with decreasing mesh size. Moreover, there is one type of error that does not automatically vanish, namely those effects that are due to the geometry of the grid. Any intrinsic regularity (such as is introduced by using a Cartesian grid) will, by Noether’s Theorem, produce spurious conservation laws. These are sometimes

innocuous, but may be rather vicious under some circumstances. One of the simplest examples is that of the Short Characteristics method [5] in radiation transport, in which the numerical diffusion along the axes is negligible compared to that along the diagonals, resulting in spiky features along the axes when modeling isotropic outflow.

It is this deficiency in particular that we wish to address here. The solution seems to be obvious: pick an irregular grid that is locally isotropic in the mean. The question then immediately arises, how to construct such a grid, and especially how to design its properties (and the corresponding algorithm that describes the physical process).

In this paper, we present the following procedure for the construction of a computational grid. First, we represent the transport medium by a point distribution. Second, these points serve as nodes in a Delaunay triangulation [6], a much used tiling of space described in Section III. Third, we specify the physical operators that act along the Delaunay lines that connect the points (in our specific example, these operators constitute a Markov transport process). The whole procedure is carried out in such a way that the point distribution (often called *point process* in the mathematics literature) optimizes the solution. By ‘representing’ we might mean a Poisson process with N points with an intensity $n_p(\vec{x})$ in 3-space, such that for $N \rightarrow \infty$ the function n_p is proportional to the matter density n . A more general case, which we will fully exploit below, is a Poisson process with $n_p = f(n)$, in which the function f is designed to optimize certain transport properties of the algorithm.

*Electronic address: ritzerveld@strw.leidenuniv.nl

We believe that this approach is extremely general and flexible, and may be used profitably in statistical physics, hydrodynamics and even quantum mechanics (cf. [7]). In this paper, however, we focus only on *linear* transport problems, of particles moving through a background medium. Thus, we ignore any nonlinear particle-particle interaction terms. For definiteness we choose to illustrate it here by means of one specific example, namely the transport of radiation through a medium that is optically active (scattering, absorption, ionization). In that case, the ‘transport medium’ consists of gas and dust (in other cases the Delaunay lines might represent, say, communication channels, or gluons connecting quarks).

B. Transport processes

The Master Equation (ME) is the mainstay of stochastic processes [8], and can be used, in its most general form, to describe the redistribution of probability in some abstract space. It relates the rate of change of some density, or distribution function, to one or more gain and loss terms which describe interactions, in the most general sense. A very successful version of the ME is the versatile Boltzmann Equation (BE), widely used in transport theory [9]. In abstract form, it reads:

$$\mathbf{D}f = \mathbf{C}f, \quad (1)$$

in which \mathbf{D} is a drift operator and \mathbf{C} a collision operator, and f is a probability density, well-defined on a certain region of this abstract space.

One can project the operators in Eq.(1) onto the phase space Γ , to obtain one particular form of the BE, that can be used to describe the flow of particles through an active medium with which they might interact. In the absence of external forces,

$$\left[\frac{\partial}{\partial t} + \vec{v} \cdot \vec{\nabla} \right] f(\vec{\mu}) = \frac{\partial f(\vec{\mu})}{\partial t} \Big|_{\text{coll}}, \quad (2)$$

in which $f(\vec{\mu}) = f(\vec{x}, \vec{v}, E, t)$ is the probability density for particles at position \vec{x} , with velocity \vec{v} , with an energy E , at time t . The collision terms on the right hand side of Eq.(2) can be written down more explicitly in terms of gain and loss terms, once the types of interactions are known. In the case of scattering one would obtain:

$$\begin{aligned} \frac{\partial f(\vec{x}, \vec{v}, E, t)}{\partial t} \Big|_{\text{coll}} &= \int_{4\pi} d\vec{v}' \Sigma(\vec{v} \cdot \vec{v}', E) f(\vec{x}, \vec{v}', E, t) \\ &\quad - f(\vec{x}, \vec{v}, E, t) \int_{4\pi} d\vec{v}' \Sigma(\vec{v} \cdot \vec{v}', E) \end{aligned} \quad (3)$$

in which $\Sigma(\vec{v} \cdot \vec{v}', E)$ is the macroscopic differential scattering cross section, equivalent to the reciprocal of the local mean free path. In general, additional particle-particle and particle-medium interactions can be added to the

collision term, each interaction having its own cross section, contributing to the total redistribution kernel.

Originally used to describe the behaviour of classical gases, the generalized BE has shown its use in a wide variety of transport processes, from that of the transport of neutrons in nuclear reactors, to that of photons in superfluids, of radiation through stellar atmospheres, to even the behavioural patterns of humans in a large crowd. As such, linear and non-linear representations of Eq.(1) have been analyzed rigorously [10], but in almost all cases it is nigh to impossible to find closed analytic solutions, even more so if one wants to find time-dependent solutions. It is therefore of the utmost scientific importance to develop reliable, fast and flexible numerical methods, which are able to deal with the complexities of transport theoretic problems.

We stress again that, in this paper, we present our method only as a tool for solving linear transport problems, such as the transport of electrons, neutrons and photons through a background medium. Extensions to the nonlinear regime are possible, but we will only hint at that in Sect. III A 3.

C. Numerical methods

Extant numerical transport methods come in a wide variety of forms, depending on which type of transport process Eq.(1) represents. In general, the computational techniques can be split into two categories: deterministic and stochastic methods. Because of the advent of ever more powerful computers, stochastic methods have been developed which mimic the behaviour of the flow itself. Monte Carlo methods based on random Markov processes are very popular, because they are conceptually simple and they are very easy to adapt to parallel processing, considering they are local. Most notably, the Direct Simulation Monte Carlo (DSMC) [11], and variants thereof, have been widely used to study the dynamics of rarefied gases where the Knudsen number (λ/L , a dimensionless parameter relating the local mean free path to the dimension of the volume) is high.

1. Structured lattices

In almost all variants of these stochastic methods, one constructs a grid, or a lattice, on which one does not only have to discretise the distribution function (and the differential operators), but also the medium properties, and thus the interaction coefficients. The transport can henceforth be solved by, for example, tracing rays (radiation transport) or transporting fluxes (hydro-solvers) from one grid cell to the next. Although a regular mesh might seem the most obvious choice, it is known to cause several problems, inherent to its structure. If the medium is highly inhomogeneous, the global resolution of a stiff rectangular mesh has to be high enough to be able

to sample the highest Fourier component of the density spectrum. This, of course, results in a high redundancy of grid cells in fairly homogeneous regions. This problem has been partly solved by the introduction of Adaptive Mesh Refinement (AMR) [12, 13], in which the grid refinement follows some criterion, such as the gradient in the density. Even refined grids, though, suffer from another problem all regular grids have: they are known to break physical symmetries. Because the cells are congruent, albeit not of the same size, the rotation group of the lattice is not isomorphic to $SO(d)$, but to some discrete subgroup thereof. Also, the fixed widths of the cells impose a non-physical constraint, and are known to break translational symmetry. Several different communities have described these shortcomings in different forms. DSMC methods require the mesh size to be much smaller than the local mean free path, in order to accurately capture the flow features. When the density of the gas locally becomes very high, the computational cost will increase tremendously, imposing a difficulty for extending DSMC methods to the continuum limit, where the Knudsen number $\lambda/L \rightarrow 0$. The Lattice Boltzmann (LB) community have shown [14] that the discrete rotation group of a rectangular lattice does not have enough symmetry to obtain Naviér Stokes-like equations in the limit, and they had to resort to hexagonal lattices in 2D and multi-speed models in 3D, in the absence of Platonic solids that at the same time have a large enough symmetry group *and* tessellate space [15]. Moreover, the regular lattices are known to break Galilean symmetry [16]. And, the fixed cell-widths impose constraint on the Reynolds number the method can resolve. In the lattice gauge community, it has been known for quite some time that regular (Wilson) lattices break Poincaré symmetry [7]. Supersymmetry closes on the Poincaré group by necessity, and therefore has difficulty being defined on regular lattices [17]. Moreover, because the lattices are invariant under translations of one or more cell-widths, or rotations of $\pi/2$ (or $\pi/3$ when one uses triangular grids) with respect to one of the axes, spurious invariants are introduced.

In summary, the choice of a regular grid which has nothing to do with the underlying physical problem results in the introduction of unphysical conserved quantities, and the breaking of several very physical symmetries.

2. Random lattices

Various different communities have independently found an answer to this mesh-related problem. Dispense with the regular grids altogether, and introduce *random* lattices, which will be described in more detail later on in this paper. In those fields of physics, where symmetries are most important, the lattices were used first. General relativity was discretised onto a simplicial lattice [18], even resulting in quantum gravity theories, and lattice

gauge theories were defined on similar random lattices [7]. Their use was hinted at in cellular automaton fluids [19], but was overlooked as a possibility to solve the dichotomy between symmetric and space-filling lattices in LB solvers (for a review on unstructured grids in LB methods, cf. [20]). Voronoi lattices were also recognized to be of use in the field of Dissipative Particle Dynamics [21, 22]. These random lattices still have a problem, though. In most cases, a Poisson point process lies at their basis, a result of which is that the average point-to-point distance is homogeneously of the order of $n_p^{-1/d}$, in which n_p is the density of points, and d is the dimension. For an inhomogeneous medium distribution this length scale does not have an immediate correspondence to the length scales of the physical problem, and, as said, impose constraints on the physical parameters (i.e. the Reynolds number for LB methods and the Knudsen for DSMC methods) to be resolved.

D. Outline

In this paper, we describe a new linear transport method which dispenses with the unphysical regular meshes, and uses as a basis random lattices that are locally isotropic in the mean. Moreover, the point distribution which defines the lattice represents the medium distribution, from which it follows that the length scales introduced (e.g. the mean Delaunay edge length) are not irrelevant, or just a step in a refinement sequence, but have true physical meaning. The method we will describe can be used more general to solve almost all MEs of type Eq.(1), and can be used as such, but in this paper we will specifically focus on linear transport processes of particles being transported through a (possibly dynamically evolving) background medium, with which they interact. Nonlinear particle-particle interactions are ignored, except through feedback from the background medium itself. Thus, the method as described in this paper can be used to model transport processes such as electron, neutron or photon transport. We will end this paper by focusing on one particular example of such particle transport, namely that of photon transport, and we will compare how an implementation of our method compares to analytical solutions. Hereafter, we show its flexibility in a simulation of the Epoch of Reionization in the early Universe, in which the first generation of stars cause a phase transition of hydrogen gas in the Universe, namely from neutral to almost fully ionized.

II. STOCHASTIC METHODS

A robust approach to numerically solving linear transport problems is the use of Monte Carlo methods, in which the solution for a macroscopic system is obtained by randomly sampling microscopic interactions. Because our new method bears many similarities to standard

stochastic methods, we proceed by first pointing out the essentials, after which we will describe our new method.

A. Monte Carlo transport methods

In very general terms, stochastic transport methods are set up to solve Eq.(1) as follows. The BE is split such that the drift term $\mathbf{D}f$ and the collision term $\mathbf{C}f$ can be treated separately, or subsequently. Hereto, one introduces a discrete time step Δt and discretises the abstract (position) space into cells using enough resolution, such that one can assume spatial homogeneity within each cell i . The collision term is solved for by using a recipe for the particle-particle or particle-medium interactions within each cell, and this term can be used to locally update $f_t(x_i)$. That new local version of $f_{t+\Delta t}(x_i)$ is then advected to the next cell via some recipe depending on the local velocity field and a possible external field.

The stochastic, or Monte Carlo, character is introduced by defining the way the collisions or interactions are solved for in each cell, at each subsequent time step. Hereto, one defines a stochastic game transforming the local state vector in each cell to a newly updated one, the statistical parameters of which can be described by the local cross section coefficients. How this is incorporated depends on the type of transport process. In rarefied gas dynamics, for example, DSMC methods solve for local collisions by defining cross sections based on the number and type of particles present in each cell. As such, the local interaction coefficients are determined by the transported particles themselves. As said, we will concentrate in this paper on linear transport problems, in which the transported particles only interact with a medium, by which the interaction coefficients are determined by the background medium properties only. Note that of course the medium itself might dynamically evolve, by which the cross section coefficients might locally change with time.

B. Particle-medium interactions

Assuming for now that the medium through which the particles are transported is static, we can elaborate on the general set up as described in the previous subsection. The standard approach to transporting neutrons, electron and photons through such a (possibly active) medium can be described as follows. First, one samples the whole domain onto a rectangular (possibly curvilinear) grid, with enough resolution to consider the medium properties (e.g. the density n_i) within each cell as constant. Given one or more sources inside or outside of the domain, one determines via some recipe the number or distribution of particles of a certain type originating from one or more grid cells, moving into certain directions. This determines our initial condition $f(\vec{x}_i, \vec{n}, t = 0)$, at each grid cell i .

One decouples the drift from the collision part of the transport equation by advecting the particles, or particle distribution $f(\vec{x}_i, \vec{n}, t = 0)$, from one grid cell i to the next i' in the direction of the particles velocity \vec{n} during some time Δt , whereafter one determines what happens to the particles depending on the medium properties within the new grid cell i' . The local medium properties are characterized by the local cross sections for the various interactions. These can be quite diverse, ranging from scattering to pure absorption. Each of this set of interactions $\{j\}$ has a unique coefficient $\alpha^j(\vec{x}_i)$, which is assumed to be constant throughout each cell.

More specifically, given a total interaction coefficient $\alpha(\vec{x}_i) = \sum_j \alpha^j(\vec{x}_i)$, one can show that the chance of particles *not* interacting with the medium within that cell is determined by the exponential distribution function

$$p(s) = \alpha(\vec{x}_i) e^{-\alpha(\vec{x}_i)s}, \quad (4)$$

in which one can define s as the length of the path of the particles from one cell to the next, or the width of one cell. One can randomly sample the probability function Eq.(4) using a Rejection Method or a Direct Inversion method [1], and determine whether or not particles will interact. What interaction will take place depends on the relative coefficients $\alpha^j(\vec{x}_i)/\alpha(\vec{x}_i)$. What happens as a result of an interaction depends on the type of interaction. When a particle is absorbed, it is subtracted from the total particle distribution, and when it is scattered, it will be redistributed along a different direction, and advected together with the surviving particles in the next time step. The moments of the distribution function in Eq.(4) are

$$\langle s^k \rangle = \int_0^\infty s^k p(s) ds = k! / \alpha^k(\vec{x}_i) = k! \lambda(\vec{x}_i)^k, \quad (5)$$

in which $\lambda(\vec{x}_i)$ is the local mean free path.

Thus, effectively, Monte Carlo methods for the linear transport of particles through a medium, move particles from one interaction to the next, along trajectories which have as an average length the local mean free path, which is determined by the medium properties within each grid cell.

III. METHOD DESCRIPTION

In the introduction, we extensively discussed the symmetry-breaking induced by the use of regular lattices, and how this can be resolved by the use of random lattices. In this section, we describe how we combine these random lattices together with the basic ideas of Monte Carlo transport methods to create a new method which solves linear transport equations on an adaptive random lattice.

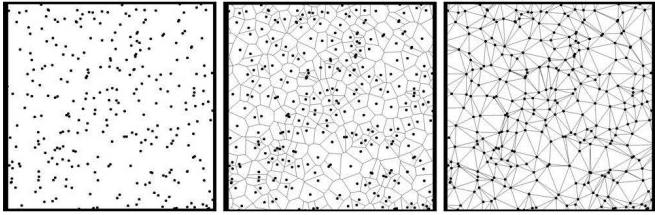


FIG. 1: *Left:* Result of a Poisson point process. *Middle:* The resultant Voronoi diagram. *Right:* The resultant Delaunay lattice.

A. Lattice construction

We will now proceed with describing how we construct our adaptive, or Lagrangian mesh, along which we will transport the particles. Note that, although we will give examples in 2D, the mesh construction method we will describe in the following is generally applicable in d -dimensional space for any $d \geq 1$.

1. Random lattices

Standard random lattices are constructed on the basis of Poisson point processes Φ , which can be defined [6] as the probability

$$\Phi = \Pr(N(A) = x) = \frac{(n_p |A|)^x e^{-n_p |A|}}{x!} \quad (6)$$

to find $x = 0, 1, 2, \dots$ points in any subset $A \subseteq S = \mathbb{R}^d$. The expectation value for Eq.(6) is $n_p |A|$, in which n_p is the point intensity, which is constant within A . This means that within every box of equal volume, the average number of points is equal. An example of this point process in 2D can be seen in Fig. 1, left. The important property of the Poisson point process is that it can be shown [6] to be translation and rotation invariant, i.e. homogeneous and isotropic. This makes it an ideal starting point to construct a lattice which retains these basic properties.

To our knowledge, the least restrictive way to construct a lattice which tessellates space, and which does not break this rotational symmetry, is the Delaunay triangulation [23]. Given a stationary point process Φ of nuclei $\{x_i\}$ in \mathbb{R}^d , the Voronoi tessellation [24] is defined as $V(\Phi) = \{C_i\}$ in which

$$C_i = \{y \in \mathbb{R}^d : \|x_i - y\| \leq \|x_j - y\| \forall x_j \neq x_i\}. \quad (7)$$

That is to say, the Voronoi cell C_i is the set of all points closer to x_i than all other points. If two Voronoi cells C_i and C_j have a common $(d-1)$ -facet (in 2D an edge, in 3D a wall, etc.), they are said to be contiguous to each other. By joining all the nuclei whose cells are contiguous, we obtain the Delaunay triangulation, which is a set of simplices and which is, by definition, dual to the

Voronoi tessellation. An example of a Voronoi and Delaunay tessellation based on a Poisson point process is depicted in Fig. 1.

There are several things to note about this random lattice. First, the recipe used to construct the Voronoi tessellation, and subsequently the Delaunay lattice, is solely based on an isotropic distance recipe (cf. Eq. (7)). This fact ensures the inheritance of rotational symmetry from the point process onto the lattice. Second, we are fortunate that specifically for Poisson-Voronoi tessellations there are several analytically derived properties available [6]. For example, the average number of edges of a Voronoi cell in 2D, and thus the expected number of Delaunay lines meeting at one vertex, is equal to 6. Thus, on average, the Poisson-Delaunay lattice is equivalent to the hexagonal lattice used in LB FHP-models [14]. In 3D, however, the number of Delaunay lines meeting at one vertex can be derived to be $(48\pi^2/35) + 2 \approx 15.54$. The fact that this number is not an even integer is a direct consequence of the fact that there is no tessellating Platonic solid in 3D which retains enough rotational symmetry, which as discussed has been known to be a problem in, for example, the LB community. Of course, every individual Voronoi cell will almost always be asymmetric, but on average it is symmetric. Thus, we have defined a random lattice which tessellates every dimension $d \geq 1$, and is rotationally and translationally symmetric, which makes it ideal for use within those parts of physics where these physical symmetries are absolutely needed. Moreover, spurious invariants associated with the breaking of these symmetries are prevented. As such, these random lattices have been used in the fluid dynamics [21, 22], lattice gauge [7], general relativity [18] and SUSY [17] communities.

2. Lagrangian random lattices

Although spurious invariant and symmetry breakings associated with rotational and translational invariance are prevented by the use of random lattices, one drawback that still remains is the introduction of an unphysical length scale, determined by the average point to point distance, or the average Delaunay line length $\langle L \rangle$. For Delaunay lattices based on a Poisson point process, the k -th order expectation value for the line length L can be derived analytically as [6]

$$\langle L^k \rangle = \zeta(k, d) n_p^{-k/d}, \quad (8)$$

in which $\zeta(k, d)$ is some geometrical constant for each pair of the value of k and the dimension d . Two often used values of this constant are

$$\zeta(1, 2) = \frac{32}{9\pi} \approx 1.132 \quad (9)$$

$$\zeta(1, 3) = \frac{1715}{2304} \left(\frac{3}{4\pi}\right)^{1/3} \Gamma\left(\frac{1}{3}\right) \approx 1.237. \quad (10)$$

The average Delaunay line length in 3D, for example, is therefore

$$\langle L \rangle = \zeta(1, 3) n_p^{-1/3} \approx 1.237 n_p^{-1/3}. \quad (11)$$

Although this line length does not have a delta function as a probability function, as in the case of a regular mesh in which the delta function peaks at a length of one cell width, but a certain spread $\sigma^2 = \langle L^2 \rangle - \langle L \rangle^2 \propto n_p^{-2/3}$, it does still have a global first order expectation value that scales with $n_p^{-1/d}$, which is a constant once the Poisson point density n_p has been chosen.

As already discussed, these fixed values cause problems when the medium distribution itself is not homogeneous, as they might underresolve high density regions, and impose constraints on parameters like the Knudsen and Reynolds numbers. More generally speaking, it causes problems in regions where the mean free path is shorter than this expectation value.

This problem has been resolved somewhat for regular meshes by introducing AMR, in which the mesh refines itself according to some predefined criterion, often based on the gradient of the pressure or density. Because here we are considering a statistical method, it is justifiably better to choose as an adaption parameter (a function of) the density of the medium. Similar efforts have been made with respect to structured grids, which have one basic congruent cell shape as a basis, in the DSMC community [25, 26], in which one wanted to make sure that the local cell sizes at least resolve the local mean free path. In this case, we do not have a regular mesh as a basis and proceed by trying to define a point process which does refine based on the local medium density.

To accomplish this, we discard the Poisson point process, which is the usual basis for random lattices. Instead, we define a point distribution, which is a convolution of a homogeneous Poisson point process Φ and a function of the possibly inhomogeneous medium density distribution $n(\vec{x})$, symbolically written as:

$$n_p(\vec{x}) = \Phi * f(n(\vec{x})). \quad (12)$$

The only constraint is the maximum number of points, or resolution, N available for the simulation. Eq.(12) amounts to nothing more than randomly sampling the function $f(n(\vec{x}))$ using a Direct Inversion of Rejection method. If the medium distribution $n(\vec{x})$ is inhomogeneous, one expects the points distribution $n_p(\vec{x})$ to be inhomogeneous too, mimicking the medium. But, as long as our number of points N is high enough, we can always zoom in far enough that locally the medium distribution is homogeneous, and the point distribution Poissonian. Thus, locally, the point distribution defined by Eq.(12) retains the rotational and translational symmetries associated with Poisson point processes.

Up until now, we have not specified the exact form of the correlation function $f(x)$ in Eq.(12). We will discuss the details hereof in the next subsection. For now, we will

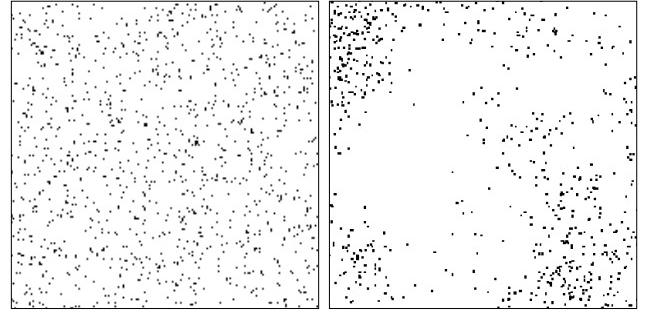


FIG. 2: *Left*: Point distribution representing a homogeneous medium; *Right*: Point distribution representing a clumpy medium.

give an example, by choosing $f(x) = x$, and plotting the resultant point distributions, for two different medium distributions. We refer to Fig. 2, in which is plotted the point distributions for a homogeneous (*left*) and an inhomogeneous (*right*) medium.

3. The correlation function

What do we choose for the correlation function $f(x)$? Because we introduced this function to ensure adaptation of the point distribution to the medium distribution, the obvious conclusion is that we need $f(x)$ to be a monotonically increasing function in x , by which the average point-to-point distances will actually be shorter in denser regions.

In light of what we discussed about Monte Carlo methods transporting particles along trajectories which have as an average length one mean free path, we can choose one particular form of $f(x)$ which makes the resultant Delaunay line lengths have a very physical meaning.

From basic transfer theory, we know that the local mean free path relates to local medium density in the following way (valid for every dimension $d \geq 1$):

$$\lambda(\vec{x}) = \frac{1}{n(\vec{x})\sigma}, \quad (13)$$

in which $\sigma = \sum_j \sigma_j$ is the total cross section, possibly consisting of many different interaction cross sections σ_j , each having its own mean free path $\lambda_j = 1/n(\vec{x})\sigma_j$. Because the mean free path is a statistical length, it scales in a different way with the medium density than the average Delaunay line length, which has an extra dimension dependence [cf. Eq.(8)]. From that, we can easily conclude that, if we choose our point distribution to sample the d -th power of the density, i.e. $f(x) \propto x^d$, or, more specifically,

$$n_p(\vec{x}) = \Phi * n^d(\vec{x}), \quad (14)$$

the length of a Delaunay line $\langle L \rangle(\vec{x})$ between two points, will scale linearly with the *local* mean free path of the

medium $\lambda(\vec{x})$ via a constant c . That is

$$\langle L \rangle(\vec{x}) = c\lambda(\vec{x}). \quad (15)$$

Thus, because we choose the point distribution to conform to the density profile of the medium according to Eq.(14), the average Delaunay line length and the mean free path have the same n^{-1} dependence, by which Eq.(15) is a *global* relation with a *global* constant c .

There are two things to note. First, the medium, and thus the medium density distribution, might evolve. In that case the lattice, which is Lagrangian by definition of Eq.(14), will evolve with it. In most relevant cases, the transport of particles through a medium is studied with respect to static media, but for several cases, such as radiation hydrodynamics, it is worthwhile to keep in mind that the medium density $n(\vec{x})$, and thus the point density $n_p(\vec{x})$, can change with time, according to what a separate hydro solver provides us with. This is also how we see a possible extension to the nonlinear transport regime of, for example, hydrodynamics. As in the DSMC method, we could define a linear Monte Carlo-like transport process through a background medium, now consisting of the particles themselves; in this case, however, the background medium, i.e. the particles themselves, can not be considered static. As such, the random lattice needs to be updated very frequently, which is a costly, and intricate task.

Second, one can choose the correlation function to be any monotonically increasing function different from the one in Eq.(14), but in that case the global constant c would change into a locally varying function $c(\vec{x})$. For example, if we choose $f(x) \propto x^e$, where $e \geq 0$, that varying function would be $c(\vec{x}) \propto n^{(d-e)/d}$.

B. Lattice properties

In the previous subsection, we described how we construct the adaptive random lattices based on the medium density distribution. Before we set out to define the way one can transport particles on this lattice, we need to discuss the exact statistical properties of the lattice, and the errors associated with its stochastic nature.

1. Distributional equivalence

The linear correlation of Eq.(15) can be taken one step further, by relating Eq.(5) to Eq.(8), and recognizing that, by choosing a correlation function $f(x) \propto x^d$, not only the first order moment, but *all* k -th moments of the exponential distribution in Eq.(4) will scale linearly with the k -th order expectation of the Delaunay line length, i.e.

$$\langle L^k \rangle(\vec{x}) = c(k)\lambda^k(\vec{x}), \quad (16)$$

in which $c(k)$ is still a global constant, but one that now depends on the order k . This similarity in statistical properties is, of course, not very surprising, because the interval between two events, or the distance between two points, in the Poisson point process in Eq.(6) has an exponential distribution $p(x) = n_p |A| e^{-n_p |A|x}$ equivalent to that of the path length between two events in Eq.(4). It is to be expected that the distribution function for the average distance, or Delaunay line lengths between two of those points follows a similar distribution, with some modifications because of the dimension of space. The mean free path statistical length is one-dimensional always, not depending on dimension, so by choosing the correlation Eq.(14) we remove the dimensional dependence of the expectation values Eq.(8), by which the path lengths of particles *and* the average lattice line lengths are distributed similarly.

2. Length sampling

We described how to construct a lattice, which is homogeneous and isotropic locally, and which adapts to the medium properties. Moreover, we showed that, by choosing a smart correlation function, the lattice becomes a direct representation of the ‘free path space’ of the particles. That is, locally, the Delaunay line lengths have the same distributional properties, i.e. the same k -th order moments, as the path lengths until interaction of the particles.

Thus, the Delaunay line originating from one point in the medium has distributional properties (k -th order moments) which all scale linearly with the distributional properties of the path lengths originating from that point. Stated differently, the variance of the Poisson process is now not associated with noise, but is exactly the variance of the free path of the particle.

Given a number of points, or resolution, N , we construct one instance of the ensemble of the point distribution, and the free paths are accurately sampled at each one of those points. This could cause inaccuracies, once the medium between two points is not locally homogeneous, by which it does not have enough sampling points. Thus, we need to impose as a sampling condition, that in the vicinity of each point (or, more accurately, within each Voronoi cell), the medium can be considered homogeneous, a condition needed in almost every numerical method. Symbolically, we need $\frac{1}{n(\vec{x})} \frac{\partial n(\vec{x})}{\partial x} < \frac{1}{\langle L \rangle}$.

This can be dealt with in two ways. First, one can increase the number of points N , by which the global parameters $c(k)$ decrease correspondingly, until the condition is satisfied. It is obvious that, when $N \rightarrow \infty$, the density field is sampled continuously, and the result is exact. Second, when the available resolution is less than needed, we can construct several instances of the same point distribution, and overlay them afterwards. This is allowed, because the exponential distribution, on which the Poisson point process is based, is memoryless, by

which the individual instances are independent. Note that the sampling condition is reached faster by choosing a correlation function as in Eq.(14).

3. Angular sampling

Another variable which has to be sampled accurately, is the number of directions a particle can propagate into at each point. As said, the average number of directions at each point is fixed (6 and 15.54, in 2D and 3D, respectively), and will not increase, when the resolution or the number of instances increases. We have to take this into account, when we design a transport algorithm, in which we want to conserve momentum at each point. Momentum conservation can be imposed, as we will show in the next subsection.

If we construct many instances of the ensemble of point distributions, the directions of the lines will differ for each instance, and, eventually, the continuous rotation group $SO(d)$ will be sampled continuously. We can, however, apply the ergodic principle, and state that it is equivalent to replace the ensemble average, by a volume average. The number of directions within a volume, containing N points, scales as $O(N)$. Thus, within a certain locally homogeneous medium, the number of points, and henceforth the number of lines has to be large enough to accurately sample the unit sphere. This amounts to nothing different to stating that a random lattice is isotropic.

Thus, we can conclude that, when the number of points within a certain small region will increase towards infinity, the number of directions will follow that trend, and the angular sampling will become infinitely precise.

C. Transport

The Lagrangian random lattice, described in the previous subsection, is a direct representation of the ‘free path space’ of the medium. As such, the stochastic element of the linear transport of particles on a fixed, deterministic grid in regular Monte Carlo methods, is moved to the simple deterministic transport of particles on a lattice, which now has the stochastic properties. Moreover, we can dispense with the regular grid or the underlying medium altogether, because all the information needed for the transport of particles has been translated to the lattice, by which the transport of particles through a medium has been translated to the transport, or percolation, of particles on the graph consisting of the lattice lines.

In the following, we will describe how we can use this lattice as a basis for transporting particles, depending on what kind of process we would like to model. As an example, we will use the two elementary linear transport processes, ballistic transport (possibly with an additional absorption term) and transport through a scattering medium, and combinations hereof.

1. Ballistic transport

We will commence with describing how we can transport particles along the lattice for a medium in which the scattering cross section is negligible. Because, in this case, we transport particles which have a predetermined momentum, it is of utmost importance to define what we do at each grid point to ensure momentum conservation, and prevent numerical diffusion.

Because our method works in such a way that the line lengths correlate linearly with the mean free paths, we assume that the homogeneously distributed medium is absorptive, and that the associated cross section σ_{abs} is the only contribution to the total cross section. Henceforth, we obtain a lattice which is similar to the Delaunay graph in Fig. 2, right.

We define one point as a source of particles, sending them out along one of the lines. The particles move along the line, until they come upon the next point. They have moved along a line, which correlates linearly with the mean free path via a constant factor c_{abs} . We can exactly evaluate the value of this factor, given Eq.(8) and Eq.(14), as

$$c_{\text{abs}} = \xi(d, N, D)\sigma_{\text{abs}}, \quad (17)$$

in which $\xi(d, N, D)$ is some constant depending on the choice of dimension d , the number of points N , and of the size of the domain D . Note that each other type of particle-medium interaction j would have a coefficient c_j determined via a relation similar to that in Eq.(17).

Ideally, one would want all $c_j = 1$, but, almost always, the resolution N will result in a coefficient larger than unity, in which case the line length is larger than the local mean free path, or smaller than unity, in which case it is the other way around. This can be taken care of by defining how the interaction, in this case absorption, is accounted for at each point.

Hereto, we define the incoming number of particles, or intensity, as I_{in} , and determine the outgoing intensity as

$$I_{\text{out}} = I_{\text{in}} e^{-c_j}, \quad (18)$$

which is equivalent to the familiar $I'(x) = I_0 e^{-x/\lambda}$, which can be derived from Eq.(4). Note that, with c_j being a global constant, e^{-c_j} is a global constant too, which can be determined a priori, via Eq.(17). In some cases, when the resolution N is high, and thus the coefficients c_j small, it might be useful to approximate $e^{-c_j} \approx (1 - c_j)$.

Locally the number of particles absorbed can be exactly evaluated as

$$I_{\text{abs}} = I_{\text{in}} (1 - e^{-c_j}), \quad (19)$$

which ensures that this method conserves particles exactly.

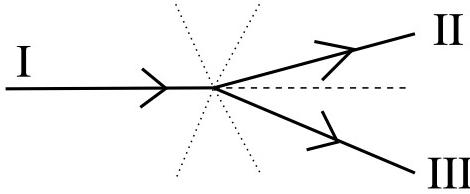


FIG. 3: Example of a number of Delaunay lines meeting at a node. Incoming particles (along line I) interact with the medium at the node, and the remaining particles should continue along the dashed line to conserve momentum. Choosing lines II and III as outgoing ensures conservation of momentum, on average.

The remaining particles I_{out} have to be sent out along one of the lines emanating from this point. The directions we choose depends on whether we want to conserve momentum, or if the interaction was with a scattering medium, in which case we want to distribute the particles isotropically. In the present case, we need to ensure momentum conservation, by which we have to choose as an outgoing line one which is in the same direction as the incoming one.

As we already discussed in the previous subsection, the lines only sample the unit sphere on average, and although the Voronoi cells are cylindrically symmetric, on average, with respect to every incoming Delaunay line, every one particular cell will deviate from that. This has as a result, that almost always, there is no outgoing line in the same direction as the incoming one. Thus, particles are deflected from their original direction by the irregularity of the grid. This can be viewed as the introduction of space-dependent inertia forces, and it is very important to keep track of these, especially in the nonlinear regime [27, 28], which we do not consider in this paper.

We can resolve this by doing the following. We refer to Fig. 3, which is an example in 2D space. In this Figure, the dashed line is the line along which momentum would be conserved. Instead, if we now choose the ‘nearest’ line II as the outgoing line, in which nearest can be defined in several ways, the most intuitive being that line for which the inner product with the dashed line is largest, it can be shown that momentum is conserved on average, which immediately follows from the fact that the Voronoi cell is axisymmetric with respect to every incoming Delaunay line. However, from a numerical implementational point of view, it is more elegant to split up the two particles and distribute one of them along line II and the other along line III, being the next to most straightforward path. More generally, if the transported quantity is a continuous entity, we have found it is most efficient to split this quantity into d equal parts, in which d is the dimension of our computational domain, and distribute each of these parts along the d most straightforward paths. Because by definition, the Poisson-Voronoi cells are axisymmetric with respect to every associated Delaunay line in every dimension $d \geq 2$, we know that a similarly modified set

of rules will ensure conservation of momentum.

To demonstrate conservation of momentum, we constructed a 2D random lattice of a Poisson point process with a relatively resolution of $N = 5 \cdot 10^4$ points, which is sufficient to show the trends and is still coarse enough to clearly show noise. Next, we define one point as source of particles, each emitted with the same momentum vector. If the particles were defined to be photons, this source could for example be called a laser beam. The rules at each site are chosen such that the two most straightforward paths with respect to this momentum vector are chosen, and that the incident package of particles are split in two and continue along these two lines. We follow the particles until they hit the absorbing boundary. The result is plotted in the bottom half of Fig. 4, in which we plotted the logarithm of the number of particles at each vertex, given that the source emits a high number of particles.

One immediately sees that, on average, the resultant momentum vector is in the original direction (by the symmetry properties of the Delaunay lattice [6]), but that there is some inherent noise associated with the use of these random lattices. Fortunately, one can prove that this noise vanishes in the limit of $N \rightarrow \infty$, by recognizing that in d -dimensional space the propagation of a particle on a lattice with this set of rules is equivalent to the process of an anisotropic random walk on a graph. An exact mathematical derivation is given in Appendix A, and we give a simplified version here.

Defining a signal $\mathcal{S} \propto \langle L \rangle n$ as the (average) distance travelled after n steps, and realizing that for an anisotropic random walk we have a noise $\mathcal{N} \propto \langle L \rangle \sqrt{n}$, we obtain a signal to noise ratio $\mathcal{S}/\mathcal{N} \propto n^{1/2}$ (which is of course similar to the famous inverse square root law of Monte Carlo methods). This is why the beam in Fig. 4 does not diverge. Given a particle location x along the momentum vector at a distance s from the source, we know that the average number of steps n for a particle to reach s scales as $n \propto s/\langle L \rangle$ in which $\langle L \rangle \propto N^{-1/d}$. From this, we can conclude that the signal to noise ratio at a distance s from the source scales as

$$\frac{\mathcal{S}}{\mathcal{N}} \propto N^{1/2d}. \quad (20)$$

Thus, in the limit of $N \rightarrow \infty$, momentum is conserved exactly, not just on average. In other words, the width of the beam will shrink to zero. Note that, when N increases, the interaction coefficients $\{c_j\}$ will decrease.

Thus, we have shown that choosing this transport recipe for ballistic transport will not only conserve momentum on average, but it does also deal with the coarse angular sampling at each point of the lattice, ensuring that any numerical diffusion will be minimal. Moreover, we have shown that this artificial widening of the beam will go to zero for infinite resolution.

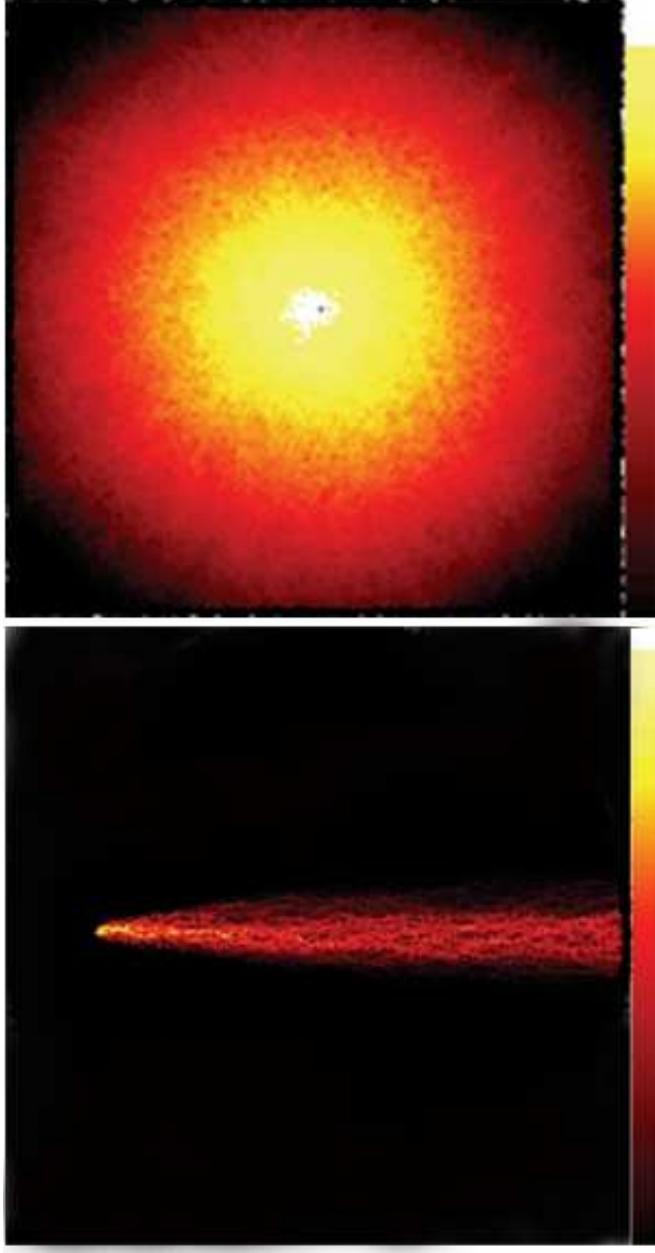


FIG. 4: (Color online) The result of two simple test on a 2D Poisson-Delaunay random lattice with $N = 5 \cdot 10^4$ points. Both are logarithmic plots of the number of particles at each site. *Top*: illustration of a scattering transport process; *Bottom*: illustration of the conservation of momentum by means of the transport of particles with constant momentum vectors.

2. Transport through scattering media

In the previous subsection we discussed, how to do transport of particles through a medium for which the interaction conserves momentum, and does not change the original momentum vectors of the particles. It is also possible that one of the interactions $\{c_j\}$ can also influence the vectorial properties of the particles, for example

when the transport is through a scattering medium.

Without loss of generality, we assume a similar homogeneous medium as before, but now, without absorptive properties, but *with* an extra interaction coefficient, σ_{scat} , which accounts for the scattering properties of the medium. In this case, the method for transportation is rather similar as in the previous subsection, in that we propagate the scattered particles $I_{\text{out}} = I_{\text{in}}e^{-c_{\text{scat}}}$ along the d most straightforward paths. What differs, is that the retained particles $I_{\text{ret}} = I_{\text{in}}(1 - e^{-c_{\text{scat}}})$ will now be redistributed isotropically. Hereto, one can choose to propagate an equal section along every one Delaunay line. Of course, the angular sampling is not perfect, but similar considerations as made in Sect. III B 3 will ensure angular resolution.

An example of one such experiment is depicted in Fig. 4, top, in which we define several point sources as a small central region as sources, and for which we let the particles propagate along the lattice according to the rules described above. We take several points as a source in this case, because this will ensure that we have many different original directions. Note that the result is very isotropic, as is to be expected of this random lattice. Moreover, the result exactly depicts the noise associated with the finite angular sampling, and thus the finite amount of points.

3. General interactions

The particle-medium interaction can consist of many different types, each one having its own contribution to the right side of Eq.(1). Each one of these interactions has its own cross section σ_j , and its own corresponding linear correlation coefficient c_j . Note, however, that the conversion factor $\xi(d, N, D)$ in Eq.(17) is the same for all interactions, given one choice of d , N and D .

The total interaction factor $C = \sum_j c_j$ determines what factor of the incoming particles interacts, in one way or the other, $I_{\text{out}} = I_{\text{in}}e^{-C}$. Every single interaction in the set $\{c_j\}$ has its own contribution $\frac{c_j}{C}e^{-C}$. When our resolution is conveniently high enough, and thus our factors $\{c_j\}$ correspondingly small, we can approximate the total factor of the particles that interact, as $(1 - e^{-C}) \approx C$, by which one can conclude that every interaction retains a factor c_j of the incoming particles. This last approximation can be implemented more efficiently code-wise.

What happens to particles, when they interact, of course depends on the type of interaction, and the type of transport process in general. Particles can be redistributed isotropically, and continue on through the medium, or they can influence the medium, by heating, ionization, or something similar. They might even be re-emitted as different particles. All these type of interactions can be implemented very easily, and the feedback of the particles on the medium is defined straightforwardly. One might even have multiple species of particles, in which case there would be a cross section σ_j for each

type of particle, and the set of factors c_j would change into a matrix, which could include the transformation of one particle into the other.

We conclude, by noting that one needs to aim at having enough resolution to ensure $c_j \leq 1$. In cases, where scattering can be neglected, the solution based on Eq.(18) would still be exact, as we could not resolve the space between the points anyway, but when wanting to incorporate scattering, one might not accurately resolve the diffusion coefficient, when $c_{\text{scat}} \gg 1$.

D. Time stepping

Up until now, we have not elaborated on what we define as a time step. As in linear Monte Carlo methods, how to define the time step depends largely on the transport problem at hand. We can mostly define two distinct cases: A) the interaction dominated limit ($\lambda/L \ll 1$), and B) the free streaming limit ($\lambda/L \geq 1$).

The first case is what is mostly encountered in linear transport problems. Here, the mean free paths for the particles are so small compared to the size of the domain, that we know the particle will be absorbed somewhere within the domain. Given that the speed of the particles is very high, one can, for each time step, let the particles that are emitted in that interval move along the lattice, moving from one interaction point to the next, until they are annihilated. The time steps do have to be chosen such that one can accurately sample time dependent source functions, or such that, when absorption is followed by ionization (cf. Sect. IV C), one can accurately follow the ionization front. Equilibrium solutions can be found, if they exist, by having a constant number of particles being emitted at each time step, that will exactly compensate for the number of particles that are absorbed, or leave the computational domain.

The second case is not very interesting from an particle-medium interaction point of view, because the densities and cross sections are of such form that no interactions are expected to occur. Thus, the particles can stream freely through quasi-vacuous space. In this case, one may be interested in following the particles' trajectories, and the time steps are then dictated by the size of the domain divided by the speed of the particles.

An interesting, albeit a bit artificial, intermediate case is the one described in Sect. III C 2. In this case of pure scattering, the particles are not expected to be annihilated. If we would not stop the particles at one point, they would eventually leave the domain. Thus, if we want to, for example, accurately follow the spherical wavefront around a point source in a homogeneous, scattering medium, as it expands with time according to the familiar Gaussian diffusion profiles, we need to keep a clock for each particle, such that it does not take more mean free path steps as is allowed by, for example, its maximum speed.

It is clear that the relevant time step criteria depend

greatly on what it is one is trying to solve for. In each case, however, one can define the relevant time stepping unambiguously. Most linear transport problems we are interested in, including the one discussed in Sect. IV C, will be of type A.

E. 3D and beyond

Because it may not seem obvious how our method is trivially extendible to \mathbb{R}^d , given all the 2D examples we gave, we will explicitly state how this extension is automatically achieved by the lattice construction procedures we gave.

First, the recipe for constructing our adaptive point process using Eq.(14) is applicable and valid in general dimensional space. Second, the procedure for constructing the Delaunay triangulation, and the corresponding recipe for the Voronoi tessellation Eq.(7), can be used and implemented efficiently in every \mathbb{R}^d [6]. Because we chose a correlation function that includes the dimension of space $n_p \propto n^d$, we made sure that the linear relation between all line lengths of the resultant lattice and the local mean free paths is valid in every dimension, cf. Eq.(16). Thus, we have obtained an adaptive random lattice that adapts to the medium in exactly the same for for every \mathbb{R}^d .

As we have already pointed out, the resultant transport process is nothing more than a walk from on interaction event to the next (as it is in Monte Carlo methods), or, alternatively, a walk from one vertex to the next along a d -dimensional Delaunay graph that has mean free path-like line lengths. The recipes for the different types of processes (ballistic, scattering, and combinations) can be extended trivially to \mathbb{R}^d . For example, for ballistic transport, we do not choose the two most straightforward paths, but the d most straightforward. The analysis concerning the conservation of momentum in the appendix is valid for every dimension, because it only assumes that the typical Voronoi cell is cylindrically symmetric around every associated Delaunay edge. This symmetry is a natural consequence of the motion invariant property of the underlying Poisson point process (see [6], for more details).

The only things that change from dimension to dimension are the average number of lines emanating from a typical point, or the number of walls of a typical Voronoi cell (6 in \mathbb{R}^2 , 15.54 in \mathbb{R}^3 , etc.), and the geometrical constant $\zeta(k, d)$ in the expressions for the k -th order moments of the Delaunay line lengths Eq.(8). These are mere constants that one has to incorporate when implementing the numerical method.

When using a regular grid, and the often associated finite-differencing, or some nontrivial form of interpolation, it is often very nontrivial to extend these operations to higher dimensional space. Because we chose to use an adaptive random lattice on which a (random) walk is performed, these difficulties are resolved, because both the

lattice construction techniques and the (random) walks are trivially defined in every \mathbb{R}^d .

F. Efficiency

It is straightforward to implement our method using one or the other programming language. We have already done this, using C++, and we will describe an application of this so-called **SimpleX** package in the next section.

For now, we will describe the basic steps of the algorithm. First of all, there have to be several pre-processing steps: 1) create a point process matching Eq.(14), in which the medium density function can be an analytic function or some data array from some other simulation; 2) construct the Delaunay triangulation; 3) determine the properties of the resultant lattice, namely the global interaction coefficients $\{c_j\}$, and the d most straightforward paths with respect to the other lines. All these values are fixed during the rest of the simulation. As an illustration, given an efficient tessellation code, these preprocessing steps can be completed within one minute wall-clock time on a simple desktop computer, for a resolution of $N = 10^6$ points.

Once the lattice and all its properties are known, the transport can commence. If we define one iteration of the algorithm as advancing particles from one point to the next, and subsequently performing the interactions at each point, and that for each point, we can make an estimate of the operation count of the algorithm.

Each point can be dealt with independently of all the others (cf. the Markov property of Monte Carlo methods). Given that the particles can be redistributed along other angular directions at each point, each line has to treated separately, but this is just a geometrical constant, given the dimension d , which does not scale with any resolution. Linear multiplications as in Eq.(18) have to be performed for each interaction j , and for each different particle species. The redistribution along the lines is just a pointer operation. Thus, the total operation count is of the form

$$N_{\text{ops}} = N_{\text{int}} N_{\text{spec}} N_p, \quad (21)$$

in which N_{int} , is the number of different interactions or non-negligible cross sections σ_j , N_{spec} is the number of species, and N_p is just the resolution, or number of points. Thus, just focusing on the resolution, the operation count of this method scales as $O(N)$.

It should be noted that this is independent of the number of sources. As it turns out, most other transport methods scale with the number of sources, by which it is extremely time-consuming to do realistic calculations for a large number of inhomogeneously distributed sources. With this method, this is now feasible, even on a simple desktop computer, as we will point out in the next section.

We end this section by pointing out that our method bears much resemblance to cellular automata methods,

in the sense that we have a global set of interaction coefficients, and a global set of rules, applied locally. Each point is influenced only by its neighbors (via the Markov criterion). Thus, like most cellular automata methods, our method can be parallelized trivially. Of course, this would be much less trivial, when nonlinear terms would be included.

G. Concluding

In this section, we have described how our new method works. We have shown how to construct a Lagrangian random lattice, which mimics the medium properties in such a way that locally all line lengths have the same distributional properties as the particles' path lengths. As such, the method can handle any geometry of the medium. The grid retains the translational and rotational symmetries, inherent to most physical problems, whilst at the same time adapting to the medium properties, by which small mean free paths and rapidly fluctuating regions of the medium will not be undersampled. The resultant lattice is a direct representation of the free path distribution space of the medium, by which the stochastic character of the the particles' trajectories as in regular Monte Carlo has been lifted to that of the grid itself, which can be shown to sample space, and all angular directions exactly when the resolution goes to infinity. Particles can be easily transported along the lattice lines, with the interactions taking place at each grid point, via an interaction coefficient c_j , which is directly proportional to the interaction cross section. We have discussed how time-stepping is introduced, depending on the linear transport problem we try to solve for, and we have demonstrated how the often not very trivial extension from \mathbb{R}^2 to \mathbb{R}^d is trivial for our method. Moreover, we have shown that the operation count of an implementation of the method is $O(N)$, which makes the method fast, even when increasing resolution, or performing 3D calculations.

IV. RADIATION TRANSPORT

As already briefly touched upon in the previous section, we have implemented the method described in this paper in a C++ package, we called **SimpleX**, named after the elementary constituent of the Delaunay lattice. For the grid construction phase, we make use of the open source package QHull [29], which has been proven to be as fast as the mathematically established limit $O(N \log N)$ for $d \leq 3$ and $O(N^{\lfloor d/2 \rfloor})$ for $d \geq 4$.

The **SimpleX** package was specifically designed to solve the equations of radiation transport. This radiative transfer comes in a wide variety of forms, depending on the (astro-)physical applications one is looking at. In the following, we will concentrate on one particular example

SimpleX has been used for, and that is the propagation of ionizing photon in the Early Universe.

A. Reionization

When the Universe had an age of about 400,000 years, its expansion had caused it to cool down such that the hydrogen recombination rate was higher than the ionization rate, by which almost all protons and electrons recombined into neutral hydrogen, making the Universe opaque to ionizing radiation. It took quite some time, before the initial density perturbations gave rise to the first generation of stars (dubbed *population III*) and quasars, which could end these Dark Ages by producing the first new supply of ionizing photons capable of reionizing the neutral hydrogen. This period, which starts with the first new photons being produced, and ends when all the hydrogen has been ionized, is believed to have happened in the redshift span $6 < z < 20$ (i.e. 150 million - one billion years after the Big Bang), and has the telltale name of *Epoch of Reionization* (EOR) [30, 31, 32]. This part of cosmology has received considerable attention in recent years, because we believe to be at the verge of observing signatures of the EOR itself, and understanding what physically happens when reionization starts is absolutely mandatory.

The initial density perturbations will enhance and enter a stage of non-linear growth, until, at the dawn of the EOR, dark and gaseous matter is distributed along a very inhomogeneous filamentary structure, in which the density range between high density filaments and low density voids can span many orders of magnitude. The photon consumption is dominated by the small scale structure, which one therefore needs to resolve, but the resultant ionization bubbles can have sizes which span sizes many magnitudes larger. It is therefore of the utmost important for methods trying to model the EOR to have large enough simulation boxes in combination with a resolution range large enough to be able to resolve the small scale structure. This, together with the inhomogeneous distribution of the matter, *and* the sources, makes it an ideal situation where the Lagrangian aspect of our method is aptly suited. Moreover, the large number of sources involved slow other methods down severely, sometimes even beyond the reach of modern supercomputers. The **SimpleX** method does not have this scaling property, and can therefore be easily used, even in these extreme cases.

B. Setup

In this specific case, the medium density $n(\vec{x})$ is determined by the hydrogen density $n_H(\vec{x})$, which is provided by dark matter simulations which let the matter distribution evolve from initial perturbations in the microwave background radiation up until the beginning of the EOR. Our code translates this density distribution into a point

distribution via Eq.(14). The standalone hydro-code also provides us with a catalogue of where the first sources have formed, together with their intensities, conveniently expressed in terms of number of ionizing photons emitted per second.

Each point is designated the same global interaction coefficient c_{ion} , associated with the photo-ionization interaction, but we include an extra factor $\chi \in [0, 1]$, which is a parameter that accounts for the feedback to the medium, and keeps track of what factor of the local medium is still neutral. This neutral fraction is updated after each iteration, and, effectively, lowers the interaction coefficient, i.e.

$$c_{\text{eff}} = \chi c_{\text{ion}}. \quad (22)$$

An extra time-dependent effect, that is needed to slow down the resultant ionization fronts, is hydrogen recombination, which can be locally incorporated by evaluating at each point, at the end of each iteration,

$$\dot{n}_{\text{rec}}(\vec{x}) = \alpha_B(1 - \chi)^2 n_H^2(\vec{x}), \quad (23)$$

in which $\dot{n}_{\text{rec}}(\vec{x})$ is the local recombination rate and α_B is the recombination coefficient. This expression effectively ignores diffuse, or scattered, radiation, which is thought to be unimportant, although recent doubt has been shed on the validity thereof [33]. Note that, if diffuse radiation turns out to actually be a real important factor, the method described in the previous section can trivially include it.

There are three things to note. First, the time-step Δt per iteration is determined by how much photons we emit per iteration, and can be made arbitrarily small, in order to accurately resolve effects like Eq.(23). Second, we can built in a limiter at each point, to make sure the velocity of the ionization front will not exceed the speed of light. And, third, in this case of radiation transport, it is extremely important that we have defined our method to be particle conserving, because precisely this ensures the ionization fronts to have the correct speed.

C. Testcase

There is one classical problem, which serves as an excellent testcase for all cosmological radiation transport code, and that is the one of an ionization front expanding in an initially neutral and uniform medium [34, 35]. It so happens that the exact analytical time-dependent solutions are known for this problem.

Given that the medium has a homogeneous density n_H , and a central monochromatic source emits photons with energy $h\nu = 13.6$ eV at a rate \dot{N}_γ , the solutions for position $r_I(t)$ and velocity $v_I(t)$ of the ionization front are

$$r_I(t) = r_S \left(1 - e^{-t/t_{\text{rec}}}\right)^{1/3} \quad (24)$$

$$v_I(t) = \frac{r_S}{3t_{\text{rec}}} \frac{e^{-t/t_{\text{rec}}}}{\left(1 - e^{-t/t_{\text{rec}}}\right)^{2/3}}, \quad (25)$$

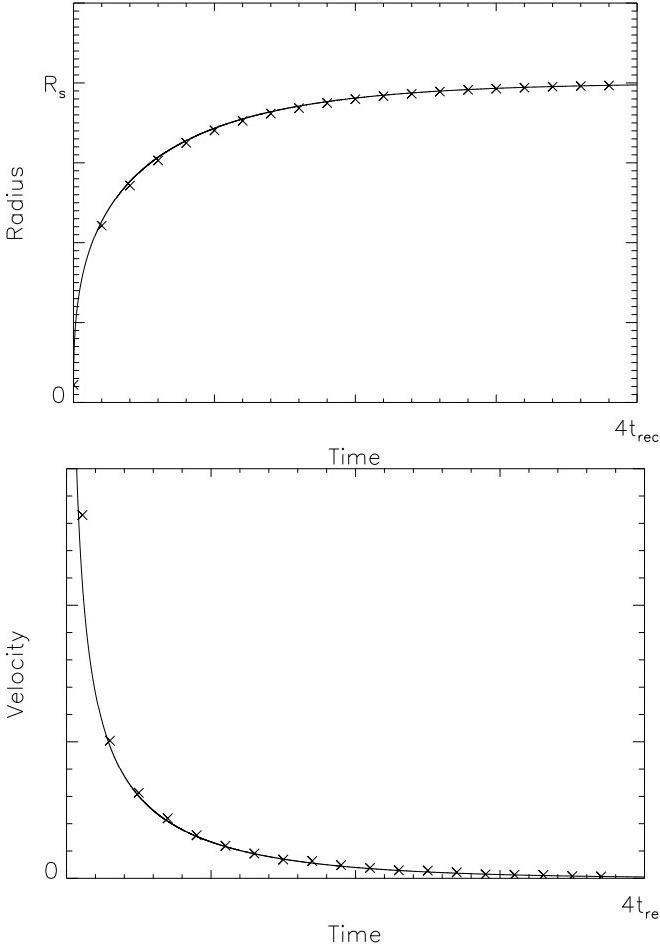


FIG. 5: Plots of the position (*top*) and velocity (*bottom*) as a function of time for an ionization front in a homogeneous medium. Crosses indicate the data produced by **SimpleX**, and the solid line represents the analytic solutions Eqs.(24) and (25).

where

$$t_{\text{rec}} = \frac{1}{\alpha_B n_H} \quad (26)$$

is the local recombination time and

$$r_S = \left(\frac{3\dot{N}_\gamma}{4\pi\alpha_B n_H^2} \right)^{1/3} \quad (27)$$

is the asymptotically reached Strömgren radius. These solutions are valid for 3D space, but similar expressions exist for general dimensional space.

Taking some typical values, $\dot{N}_\gamma = 5 \cdot 10^{48} \text{ s}^{-1}$, $n_H = 10^{-3} \text{ cm}^{-3}$, $\alpha_B = 2.59 \cdot 10^{-13} \text{ cm}^3 \text{s}^{-1}$ and a chosen simulation time $T = 4t_{\text{rec}}$, we can run **SimpleX** with these parameters, and determine what the position and velocity of the ionization front is after each iteration. These results we can immediately compare with the analytical solutions Eqs.(24) and (25), which has been done in Fig. 5.

It is straightforward to see that the results of the **SimpleX** method agree very well with what is analytically expected. A similarly good agreement can be seen, when running the simulation for a medium with an r^{-1} or an r^{-2} density distribution, in which r is the distance to the central source. Analytic solutions are available too for these cases [36].

Unfortunately, except for these rare cases, there are no analytical solutions available for problems where the medium and the source distribution is more inhomogeneous, and this is precisely the regime in which cosmological radiative transport methods have to perform well. To provide a more quantitative basis on which one can validate the performance of ones code, a cosmological radiative transfer code comparison project was initiated last year, in which 11 different codes, of which **SimpleX** was one, were compared. The comparison consists of several testcases, one of which was described in this subsection, but also several which do not have an analytical solution. The results of this comparison project is described in the forthcoming paper [37].

D. Large scale simulation

As an illustration of what is possible with this method, we give a more realistic example of the reionization of a box which has as a size 100 Mpc comoving, in which a certain inhomogeneous matter distribution, extracted from a dark matter simulation with the code PMFAST [38], has to be ionized by a certain number of sources. A slice through the box is depicted in Fig. 6, in which the logarithm of the density is plotted. One can clearly see the filamentary structure and the high contrast between low and high density regions.

We use 10^6 points to sample this medium distribution, and via a friends-of-friends halo finding algorithm, we can designate several points as sources of ionizing photons. A total 2000 points are marked as sources and we assume that the medium is static, and that there are no radiation hydrodynamical feedback effects.

The result of the **SimpleX** reionization simulation of this medium distribution is depicted in Fig. 7, in which a volume rendering of the ionization structure of the medium is plotted for 6 different points in time. White is the still neutral and opaque hydrogen, and blue is what is already ionized, and transparent. A cutout has been made to more clearly show the inner structure of the simulated volume.

One can clearly see ionization bubbles, similar to the ones described by Eqs.(24) and (25), being blown around each source, which eventually overlap and fill the whole box. At this point the EOR ends, and, again, the Universe is transparent to ionizing radiation.

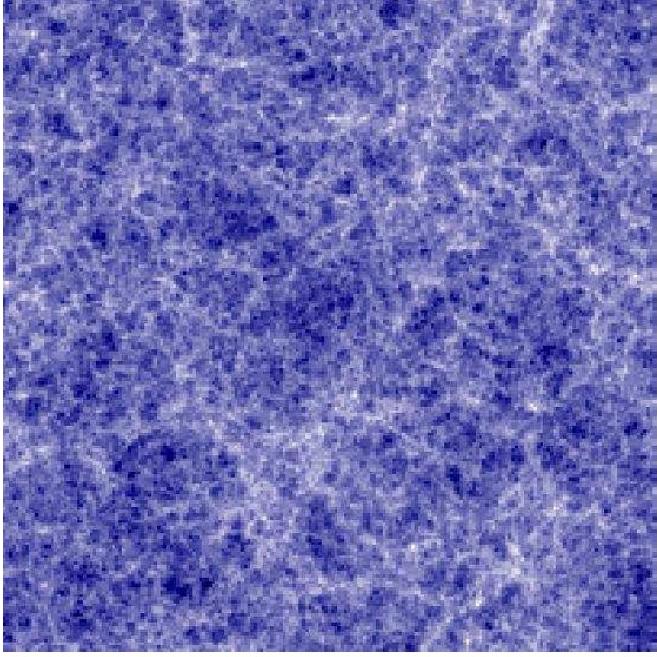


FIG. 6: (Color online) A slice through the simulation box, depicting the density distribution of neutral hydrogen on a logarithmic scale. The simulation was done using **PMFAST** [38]. This is a direct representation of the large scale structure of the Universe.

V. SUMMARY AND DISCUSSION

The simulation of the reionization of the large scale structure of the Universe, described in the previous section, runs on a simple desktop computer, finishing within an hour. Other cosmological radiative transfer methods would need weeks, even months of supercomputer time to accomplish the same task. This efficiency is due to the use of the Lagrangian lattice, which puts the usually limited resolution at places where it is needed, whilst at the same time circumventing the severe problems associated with structured grids that introduce unphysical conserved quantities. The resultant method does not scale with the number of sources, which appears to be one of the major bottlenecks of existing linear transport methods. As briefly mentioned in Sect. III A 3, it is possible to include nonlinear terms into our transport process, by discarding the static background medium and letting the random lattice evolve according to the particle density. This paper deals with linear problems only, which is why we refrain from discussing this possible extension here.

In a more general setting, we described in this paper a method which is, in principle, suited to solve linear transport equations in the abstract form of Eq.(1). Hereto, we assume that we can treat the advection and interaction parts disjointly, or subsequently. The method is new in the sense that it dispenses with the regular, structured grids altogether, and introduces a very physical one.

Focusing on transport in which the only interaction terms involve that between particles and a static medium, we extensively described how to create an adaptive grid, based on a random point process, which locally retains important physical symmetries, namely rotational and translational invariance. Moreover, the resultant lattice can be constructed in such a way that all line lengths correlate linearly with all statistical moments of the path lengths of the particles. As such, the resultant lattice is a direct representation of the free path space of the medium, which characterizes its interaction properties. The exact value of the coupling constant σ_j of each interaction is converted into a global interaction coefficient c_j , which is a quantitative measure of how many mean free paths fit into one lattice line length. After that, we showed that, in the limit of an infinite number of points, the lattice continuously samples all possible free paths and all possible angular directions.

Henceforth, we described how to do the actual transport, by solving for the advection part by moving the particles along the lattice, or graph, and solving for the interaction part, by evaluating Eq.(18) at each point. Redistribution of the surviving particles can be chosen to conserve momentum, and redistribution of scattered particles can be chosen to be isotropic. Of course, it is trivial to incorporate intermediate cases, in which the scattering is anisotropic.

We showed the results of implementing the method as a radiative transfer code **SimpleX**, with which we evaluated the classical problem of an ionization front expanding in a homogeneous medium towards the asymptotically reached Strömgren sphere. The results are in exact agreement with the analytical solutions, which is also the case for r^{-1} and r^{-2} medium distributions. The method has been compared to other codes using testcases for which there are no analytical solutions. As an illustration of what the method is capable of, we finished by using our **SimpleX** method to model the photo-ionization of the large scale structure of the Universe during the Epoch of Reionization.

From an implementation point of view, the method is very easy to use, given that Delaunay and Voronoi tessellation are widely used in many areas of science, including computer visualization, by which fast, robust and open source codes for performing the lattice construction are plentiful. Given the lattice, there is not much more to implement, apart from moving particles along a list of pointers. The expensive finite-differencing, and such, associated with the use of structured grids, was dispensed with together with these lattices. Moreover, because our interaction is implemented as a set of global interaction coefficients, with global interaction rules, each cell can be considered to be a cellular automaton. Thus, our method can be trivially parallelized.

We would like to conclude by noting that although in this paper we focused on linear transport of photon-like particles through a static medium, the method can be used much more generally. The space in which the path

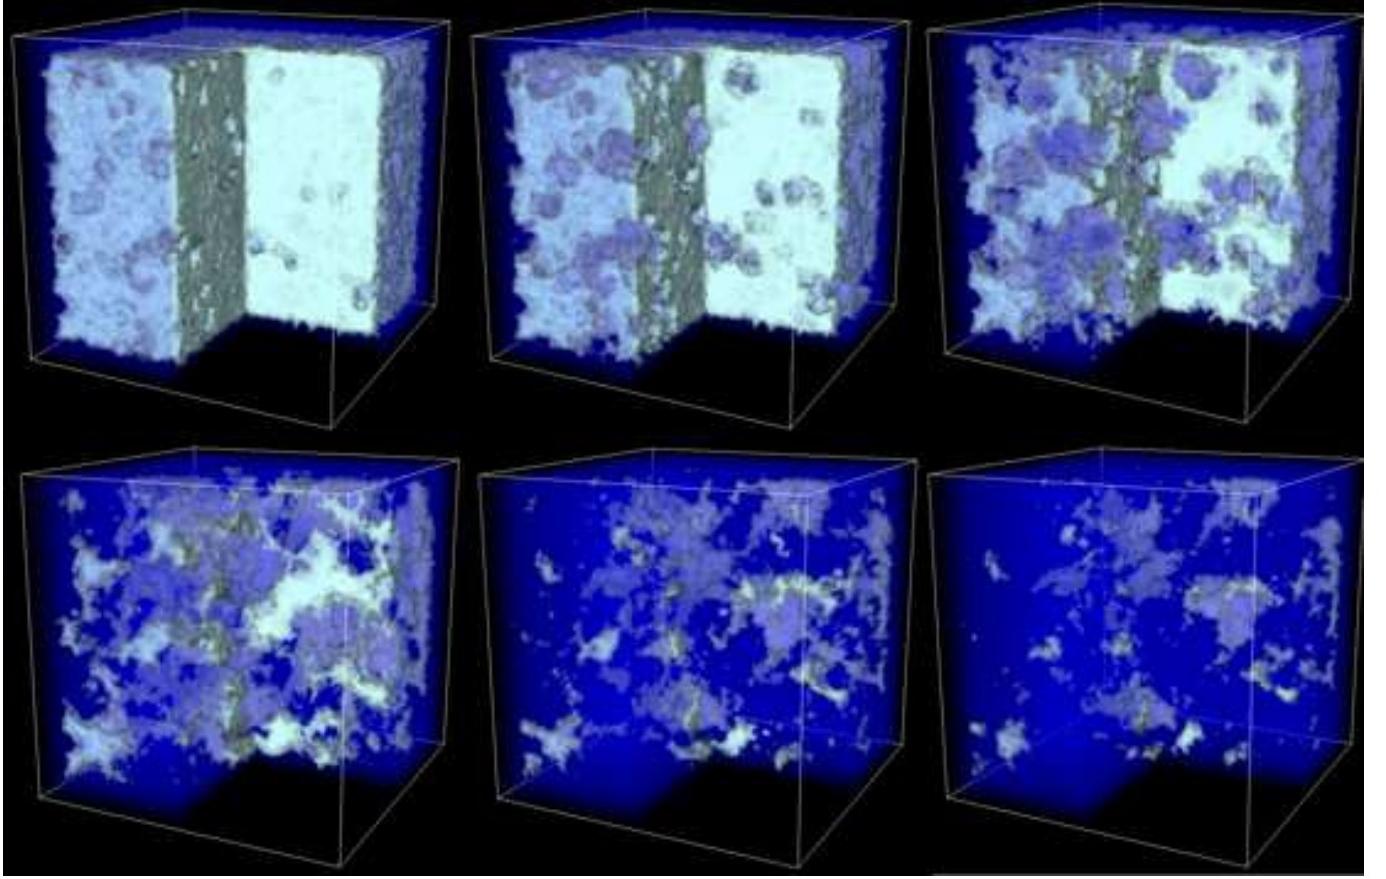


FIG. 7: (Color online) A volume rendering of the result of using the **SimpleX** method to transport the ionizing photons through the medium distribution of the box depicted in Fig. 6. The results of 6 different points in time are plotted, in which the white corresponds to the hydrogen that is still neutral (opaque), and the blue to the already ionized hydrogen (transparent).

lengths are defined does need to be defined as a subset of phase space, but can be much more abstract, and diverse. In every case, the method will consist of constructing a Lagrangian random lattice, as a direct representation of what the transported quantity would encounter, when it is transported through that abstract space.

I like to thank Yuri Levin for carefully reading the manuscript, and Garrelt Mellema, Erik-Jan Rijkhorst, and everyone participating in the comparison project (<http://www.mpa-garching.mpg.de/tsu3/>) for fruitful discussions. This work was done with financial support from the Netherlands Organization for Scientific Research (NWO), under grant number 635.000.009.

Acknowledgments

We are much obliged to Ilian Iliev for making available the dark matter simulation data. JR would

-
- [1] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge University Press, New York, NY, USA, 1992).
 - [2] R. D. Richtmyer and K. W. Morton, *Difference Methods for Initial Value Problems* (Interscience Publishers, New York, 1967).
 - [3] B. Van Leer, Ph.D. thesis, Leiden (1970).
 - [4] V. Icke, *A&A* **202**, 177 (1988).
 - [5] P. B. Kunasz and L. Auer, *JQSRT* **39**, 67 (1988).
 - [6] A. Okabe, B. Boots, K. Sugihara, and S. Chiu, *Spatial Tessellations, Concepts and Applications of Voronoi Diagrams* (John Wiley & Sons, 1999), 2nd ed.
 - [7] N. H. Christ, R. Friedberg, and T. D. Lee, *Nuclear Physics B* **202**, 89 (1982).
 - [8] N. Van Kampen, *Stochastic Processes in Physics and*

- Chemistry* (North Holland, 1981).
- [9] C. Cercignani, *The Boltzmann Equation and Its Applications* (Springer-Verslag, 1988).
 - [10] N. Bellomo, ed., *Lecture Notes on the Mathematical Theory of the Boltzmann Equation* (World Scientific, 1995).
 - [11] G. Bird, *Molecular Gas Dynamics and the Direct Simulation of Gas Flows* (Clarendon, 1994).
 - [12] M. J. Berger and J. Oliger, *J. Comput. Phys.* **53**, 484 (1984).
 - [13] M. J. Berger and P. Colella, *J. Comput. Phys.* **82**, 64 (1989).
 - [14] U. Frisch, B. Hasslacher, and Y. Pomeau, *Phys. Rev. Lett.* **56**, 1505 (1986).
 - [15] D. W. Gladrow, *Lattice Gas and Lattice Boltzmann Methods* (Springer-Verslag, Berlin, 2000).
 - [16] S. Succi, *The Lattice Boltzmann Equation for Fluid Dynamics and Beyond* (Oxford University Press, Oxford, 2001).
 - [17] M. Kaku, *Phys. Rev. Lett.* **50**, 1893 (1983).
 - [18] T. Regge, *Nuovo Cimento A* **19**, 558 (1961).
 - [19] S. Wolfram, *J. Stat. Phys.* **45**, 471 (1986).
 - [20] S. Ubertini and S. Succi, *Progress in Computational Fluid Dynamics* **5**, 85 (2005).
 - [21] P. Espa  nol, *Europhysics Letters* **39**, 605 (1997).
 - [22] E. G. Flekk  y, P. V. Coveney, and G. de Fabritiis, *Phys. Rev. E* **62**, 2140 (2000).
 - [23] B. Delaunay, *Classe des Sciences Math  matiques et Naturelles* **7**, 793 (1934).
 - [24] G. Voronoi, *Journal f  r die Reine und Angewandte Mathematik* **134**, 198 (1908).
 - [25] A. L. Garcia, J. B. Bell, W. Y. Crutchfield, and B. J. Alder, *J. Comput. Phys.* **154**, 134 (1999).
 - [26] J.-S. Wu, K.-C. Tseng, and C.-H. Kuo, *Int. J. Numer. Meth. Fluids* **38**, 351 (2002).
 - [27] R. Benzi, S. Succi, and M. Vergassola, *Physics Reports* **222**, 145 (1992).
 - [28] I. V. Karlin, S. Succi, and S. Orszag, *Phys. Rev. Lett.* **82**, 5245 (1999).
 - [29] C. Bradford Barber, D. P. Dobkin, and H. Huhdanpaa, *ACM Trans. on Mathematical Software* **22**, 469 (1995), URL <http://www.qhull.org>.
 - [30] A. Loeb and R. Barkana, *ARA&A* **39**, 19 (2001).
 - [31] V. Bromm and R. B. Larson, *ARA&A* **42**, 79 (2004).
 - [32] B. Ciardi and A. Ferrara, *Space Science Reviews* **116**, 625 (2005).
 - [33] J. Ritzerveld, *A&A* **439**, L23 (2005).
 - [34] B. Str  mgren, *ApJ* **89**, 526 (1939).
 - [35] L. Spitzer, *Physical processes in the interstellar medium* (New York Wiley-Interscience, 1978. 333 p., 1978).
 - [36] G. Mellema, I. T. Iliev, M. A. Alvarez, and P. R. Shapiro, *New Astronomy* **11**, 374 (2006).
 - [37] I. T. Iliev, B. Ciardi, M. A. Alvarez, A. Ferrara, N. Y. Gnedin, A. Maselli, G. Mellema, T. Nakamoto, M. L. Norman, A. O. Razoumov, et al., ArXiv Astrophysics e-prints (2006), astro-ph/0603199.
 - [38] H. Merz, U.-L. Pen, and H. Trac, *New Astronomy* **10**, 393 (2005), astro-ph/0402443.

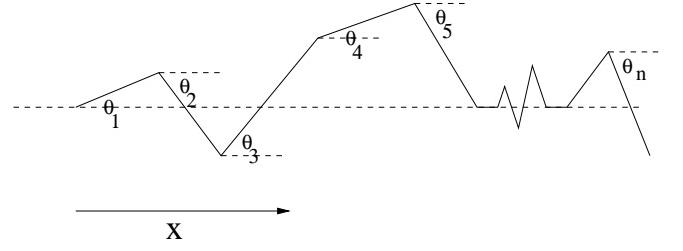


FIG. 8: One possible path of particles performing a walk of n steps on the Delaunay graph. The i -th step is parametrized by an angle θ_i , with respect to the original direction \mathbf{x} .

APPENDIX A: CONSERVATION OF MOMENTUM

The symmetry of each typical Voronoi cell [6] ensures that momentum is conserved on average, given the ballistic transport recipe in Sect. III C 1. What is more interesting, though, from an implementation point of view, is the width of the distribution around the original direction. That is, how wide would an otherwise infinitely thin laser beam become as function of the distance to the source. For an exact mathematical analysis hereof for the Poissonian random lattice, we proceed as follows.

An example of a path of a particle performing a walk in two dimensions is given in Fig. 8. The following analysis, however, will be valid in d -dimensional space. Because of cylindrical symmetry around the original direction \mathbf{x} (the distribution function of the deflection angle is symmetric), we can parametrize the i -th step of the particle's walk by only one angle θ_i , which is the angle between the i -th Delaunay edge and the original direction \mathbf{x} . Thus, the expectation value of the total displacement $\mathbf{R}_n = \mathbf{r}_1 + \dots + \mathbf{r}_n$ is

$$\begin{aligned} \langle \mathbf{R}_n \rangle &= \langle \mathbf{r}_1 \rangle + \dots + \langle \mathbf{r}_n \rangle \\ &= n \langle L \rangle \langle \cos \theta \rangle \frac{\mathbf{x}}{|\mathbf{x}|} \\ &= n \langle L \rangle \chi \frac{\mathbf{x}}{|\mathbf{x}|}, \end{aligned} \quad (\text{A1})$$

in which $\langle L \rangle$ is the average Delaunay line length, defined in Eq.(8), and

$$\chi = \int_{-\pi}^{\pi} h(\theta) \cos \theta d\theta. \quad (\text{A2})$$

Here, we have used $h(\theta)$ as a certain symmetric function, which characterises the probability distribution of the angle θ and which, in most cases, cannot be evaluated analytically. Several Monte Carlo experiments for this angle have been done [6].

The second-order expectation value can be evaluated as follows:

$$\begin{aligned} \langle \mathbf{R}_n^2 \rangle &= \langle \mathbf{r}_1^2 \rangle + \langle \mathbf{r}_1 \cdot \mathbf{r}_2 \rangle + \dots + \langle \mathbf{r}_n^2 \rangle \\ &= \langle L^2 \rangle [n + n(n-1) \langle \cos(\theta_i + \theta_j) \rangle], \end{aligned} \quad (\text{A3})$$

in which we may choose i and j randomly from the set $\{1, \dots, n\}$, as long as $i \neq j$, because the distribution function $h(\theta)$ has the same form for each angle θ_i . Using the cosine addition formula, we can reduce Eq.(A3) to

$$\langle \mathbf{R}_n^2 \rangle = (n + n(n-1)\chi^2) \langle L^2 \rangle. \quad (\text{A4})$$

Thus, the variance of the displacement is

$$\sigma_{\mathbf{R}_n}^2 = n \langle L^2 \rangle (1 - \chi^2). \quad (\text{A5})$$

When $h(\theta) \propto \delta(\theta)$, then $\chi = 1$, by which $\langle \mathbf{R}_n \rangle = \langle L \rangle n \mathbf{x} / |\mathbf{x}|$ and $\sigma_{\mathbf{R}_n}^2 = 0$ as should be expected. The exact form of a distribution function like $h(\theta)$ can probably not be evaluated, even for this well-studied Poisson case, but we can use a step function as an approximation. Thus, given that in \mathbb{R}^2 the average number of Delaunay lines meeting at a grid point is 6, we use as a step function $h(\theta) = 3/\pi$ on the domain $\theta \in [-\pi/6, \pi/6]$. This results in $\chi = 3/\pi$, by which

$$\langle \mathbf{R}_n \rangle = \frac{3n \langle L \rangle}{\pi} \frac{\mathbf{x}}{|\mathbf{x}|}, \quad (\text{A6})$$

which is very close (difference of less than 5%) to the distance along a straight line, which would be $n \langle L \rangle$. We can always, of course, rescale the lengths so as to make sure that the distance traversed equals the exact physical one.

More importantly, the variance in the displacement, in this case, is

$$\sigma_{\mathbf{R}_n}^2 = \frac{\pi^2 - 9}{\pi^2} \langle L^2 \rangle n. \quad (\text{A7})$$

We know that the results of using a step-function as distribution function gives upper bounds on the values of

Eqs.(A1) and (A5), because, similar to the deflection angle distribution function, the actual distribution function would peak around $\theta = 0$ and would decrease as $|\theta|$ increases, so we expect the actual value of $\sigma_{\mathbf{R}_n}^2$ to be smaller. Thus, we can simulate a straight line trajectory with this method, because $\mathbf{R}_n \propto n \mathbf{x}$, but with a standard deviation that increases with \sqrt{n} .

A crucial aspect is the behavior of the standard deviation, when the number of grid points N increases. Let us therefore examine a line segment in the simulation domain of length \mathcal{L} ($\leq \sqrt{d}$, if we have a $[0.0 : 1.0]^d$ domain). Because the point distribution is homogeneous, we can conclude that the number of steps to cover the line is

$$n = \xi N^{1/d}, \quad (\text{A8})$$

in which $\xi \leq \frac{\pi}{3} \zeta(1, 2) \sqrt{d}$, which can be found by using the upper bound Eq.(A6) and the Eq.(8) for the length $\langle L \rangle$ of a Delaunay line. If we combine Eq.(A7) with Eq.(A8), again using Eq.(8), we obtain

$$\sigma \propto \langle L \rangle \sqrt{n} \propto N^{-1/2d}. \quad (\text{A9})$$

Thus, we can conclude that the amount of widening of the beam will go to zero, when we increase the amount of grid points N .

Even if we do not have a large amount of points to suppress the widening of the beam, we have another effect which compensates for the widening. Namely, at each intersection the number of particles are split up into d parts. This means that the particle number at points farther away from the straight line trajectory is much less than at points close by, simply because of the fact that more paths cross each other at points close to the line.